

06-09-00

EE318062105US
June 8, 2000

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Docket No. AUS9-2000-0220-US1

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

Transmitted herewith for filing is the patent application of Inventor(s):
RICHARD LOUIS ARNDTFor: **LOGICAL PARTITIONING VIA HYPERVISOR MEDIATED ADDRESS
TRANSLATION**

Enclosed are also:

- ☒ 18 Pages of Specification including an Abstract
☒ 6 Pages of Claims
☒ 4 Sheet(s) of Drawings
☒ A Declaration and Power of Attorney
☒ Form PTO 1595 and assignment of the invention to IBM Corporation

CLAIMS AS FILED

FOR	Number Filed	Number Extra	Rate	Basic Fee (\$690)
Total Claims	19 -20 =	0	X \$ 18	= \$ 0.00
Independent Claims	4 -3 =	1	X \$ 78	= \$ 78.00
Multiple Dependent Claims	0		X \$ 260	= \$ 0.00
Total Filing Fee				= \$ 768.00

- ☒ Please charge \$ 768.00 to IBM Corporation, Deposit Account No. 09-0447.
☒ The Commissioner is hereby authorized to charge payment of the following fees associated with the communication or credit any over payment to IBM Corporation, Deposit Account No. 09-0447. A duplicate copy of this sheet is enclosed.
☒ Any additional filing fees required under 37CFR § 1.16.
☒ Any patent application processing fees under 37CFR § 1.17.

Respectfully,

Mark E. McBurney

Reg. No. 33,114

Intellectual Property Law Dept.

IBM Corporation

11400 Burnet Road 4054

Austin, Texas 75758

Telephone: (512) 823-1003

00/80/90
JC834 U.S. PTO
09/589795JC834 U.S. PTO
09/589795
06/08/00

Docket No. AUS9-2000-0220-US1

**LOGICAL PARTITIONING VIA HYPERVISOR MEDIATED ADDRESS
TRANSLATION**

5

BACKGROUND OF THE INVENTION

1. Technical Field:

The present invention relates generally to the field of computer architecture and, more specifically, to
10 methods and systems for managing resources among multiple operating system images within a logically partitioned data processing system.

2. Description of Related Art:

15 A logical partitioning option (LPAR) within a data processing system (platform) allows multiple copies of a single operating system (OS) or multiple heterogeneous operating systems to be simultaneously run on a single data processing system platform. A partition, within
20 which an operating system image runs, is assigned a non-overlapping sub-set of the platform's resources. These platform allocable resources include one or more architecturally distinct processors with their interrupt management area, regions of system memory, and I/O
25 adapter bus slots. The partition's resources are represented by its own open firmware device tree to the OS image.

Each distinct OS or image of an OS running within the platform are protected from each other, such that
30 software errors on one logical partition cannot affect the correct operation of any of the other partitions. This is provided by allocating a disjoint set of platform

Docket No. AUS9-2000-0220-US1

resources to be directly managed by each OS image and by providing mechanisms for ensuring that the various images can not control any resources that have not been allocated to it. Furthermore, software errors in the control of an OS's allocated resources are prevented from affecting the resources of any other image. Thus, each image of the OS (or each different OS) directly controls a distinct set of allocable resources within the platform.

10 A significant problem in an LPAR computer system is the mechanism that performs the isolation of partition resources from each other.

In some previous implementations, such as, for example, the IBM S/390, a product of the International Business Machines Corporation of Armonk, New York, the platform included extra hardware in each processor that added a fixed offset to all processor generated addressed. The added hardware also checked to make sure that the resulting address did not exceed a specific upper value. This extra hardware is made inaccessible to the logical partition's OSs. However, this method forces all the logical partition's directly accessible resources to be configured in a single (or at most a few) contiguous regions of the platform's address space.

20 Other partition resources, such as, for example, sparsely mapped memory and mapped input/output (I/O) devices, must be accessed indirectly through platform firmware.

Another popular method, employed by Sun Micro-Systems, Inc., as well as Intel based systems, is to electrically disconnect, in some fashion, the bus between sub-sets to be configured with the complete set

of resources that the partition needs. Thus, changing a partition's resource configuration often involves physically moving hardware between the sub-sets.

Therefore, it is desirable to have a method, system, and apparatus that isolates partition resources from each other while allowing fine grain allocation of resources to partitions without necessitating the physical movement of the hardware during reconfiguration.

Docket No. AUS9-2000-0220-US1

SUMMARY OF THE INVENTION

5 The present invention provides a method, system, and
apparatus for mediating address translation in a
logically partitioned data processing system. In one
embodiment, a firmware component receives from an
operating system within a logical partition a request to
access a physical resource. The firmware component,
10 responsive to a determination that the physical resource
has been allocated to the logical partition, modifies an
address translation table, if necessary, to allow access
to the physical resource by the operating system. The
operating system is prevented from directly modifying the
15 address translation table, thus preventing potential
interference between operating systems within the
logically partitioned data processing system.

Docket No. AUS9-2000-0220-US1

BRIEF DESCRIPTION OF THE DRAWINGS

5 The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives and advantages thereof, will best be understood by reference to the following detailed
10 description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 depicts a block diagram of a data processing system, which may be implemented as a logically partitioned server, in accordance with the
15 present invention;

Figure 2 depicts a block diagram of a logically partitioned platform in which the present invention may be implemented;

Figure 3 depicts a block diagram illustrating an
20 exemplary system for mediating virtual address translation in accordance with the present invention; and

Figure 4 depicts a flowchart illustrating an exemplary process in a firmware component for mediating address translation between various operating systems and
25 physical resources in a logically partitioned data processing system in accordance with the present invention.

Docket No. AUS9-2000-0220-US1

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

5 With reference now to the figures, and in particular with reference to **Figure 1**, a block diagram of a data processing system, which may be implemented as a logically partitioned data processing system, is depicted in accordance with the present invention. Data
10 processing system **100** may be a symmetric multiprocessor (SMP) system including a plurality of processors **101**, **102**, **103**, and **104** connected to system bus **106**. For example, data processing system **100** may be an IBM RS/6000, a product of International Business Machines
15 Corporation in Armonk, New York. Alternatively, a single processor system may be employed. Also connected to system bus **106** is memory controller/cache **108**, which provides an interface to a plurality of local memories **160-163**. I/O bus bridge **110** is connected to system bus
20 **106** and provides an interface to I/O bus **112**. Memory controller/cache **108** and I/O bus bridge **110** may be integrated as depicted.

 Data processing system **100** is a logically partitioned data processing system. Thus, data
25 processing system **100** may have multiple heterogeneous operating systems (or multiple instances of a single operating system) running simultaneously. Each of these multiple operating systems may have any number of software programs executing within in it. Data
30 processing system **100** is logically partitioned such that different I/O adapters **120-121**, **128-129**, **136-137**, and

Docket No. AUS9-2000-0220-US1

146-147 may be assigned to different logical partitions.

Thus, for example, suppose data processing system **100** is divided into three logical partitions, P1, P2, and P3. Each of I/O adapters **120-121**, **128-129**, and **136-137**,
5 each of processors **101-104**, and each of local memories **160-164** is assigned to one of the three partitions. For example, processor **101**, memory **160**, and I/O adapters **120**, **128**, and **129** may be assigned to logical partition P1; processors **102-103**, memory **161**, and I/O adapters **121** and
10 **137** may be assigned to partition P2; and processor **104**, memories **162-163**, and I/O adapters **136** and **146-147** may be assigned to logical partition P3.

Each operating system executing within data processing system **100** is assigned to a different logical
15 partition. Thus, each operating system executing within data processing system **100** may access only those I/O units that are within its logical partition. Thus, for example, one instance of the Advanced Interactive Executive (AIX) operating system may be executing within
20 partition P1, a second instance (image) of the AIX operating system may be executing within partition P2, and a Windows 2000™ operating system may be operating within logical partition P1. Windows 2000 is a product and trademark of Microsoft Corporation of Redmond,
25 Washington.

Peripheral component interconnect (PCI) Host bridge **114** connected to I/O bus **112** provides an interface to PCI local bus **115**. A number of Terminal Bridges **116-117** may be connected to PCI bus **115**. Typical PCI bus
30 implementations will support four Terminal Bridges for

Docket No. AUS9-2000-0220-US1

providing expansion slots or add-in connectors. Each of Terminal Bridges **116-117** is connected to a PCI I/O Adapter **120-121** through a PCI Bus **118-119**. Each I/O Adapter **120-121** provides an interface between data

5 processing system **100** and input/output devices such as, for example, other network computers, which are clients to server **100**. Only a single I/O adapter **120-121** may be connected to each terminal bridge **116-117**. Each of

10 terminal bridges **116-117** is configured to prevent the propagation of errors up into the PCI Host Bridge **114** and into higher levels of data processing system **100**. By doing so, an error received by any of terminal bridges **116-117** is isolated from the shared buses **115** and **112** of the other I/O adapters **121**, **128-129**, and **136-137** that may

15 be in different partitions. Therefore, an error occurring within an I/O device in one partition is not "seen" by the operating system of another partition. Thus, the integrity of the operating system in one partition is not effected by an error occurring in

20 another logical partition. Without such isolation of errors, an error occurring within an I/O device of one partition may cause the operating systems or application programs of another partition to cease to operate or to cease to operate correctly.

25 Additional PCI host bridges **122**, **130**, and **140** provide interfaces for additional PCI buses **123**, **131**, and **141**. Each of additional PCI buses **123**, **131**, and **141** are connected to a plurality of terminal bridges **124-125**, **132-133**, and **142-143** which are each connected to a PCI

30 I/O adapter **128-129**, **136-137**, and **146-147** by a PCI bus

Docket No. AUS9-2000-0220-US1

126-127, 134-135, and 144-145. Thus, additional I/O devices, such as, for example, modems or network adapters may be supported through each of PCI I/O adapters 128-129, 136-137, and 146-147. In this manner, server 5 100 allows connections to multiple network computers. A memory mapped graphics adapter 148 and hard disk 150 may also be connected to I/O bus 112 as depicted, either directly or indirectly. Hard disk 150 may be logically partitioned between various partitions without the need 10 for additional hard disks. However, additional hard disks may be utilized if desired.

Those of ordinary skill in the art will appreciate that the hardware depicted in **Figure 1** may vary. For example, other peripheral devices, such as optical disk 15 drives and the like, also may be used in addition to or in place of the hardware depicted. The depicted example is not meant to imply architectural limitations with respect to the present invention.

With reference now to **Figure 2**, a block diagram of 20 an exemplary logically partitioned platform is depicted in which the present invention may be implemented. The hardware in logically partitioned platform 200 may be implemented as, for example, data processing system 100 in **Figure 1**. Logically partitioned platform 200 includes 25 partitioned hardware 230, virtual address translation hardware 280, hypervisor 210, and operating systems 202-208. Operating systems 202-208 may be multiple copies of a single operating system or multiple heterogeneous operating systems simultaneously run on 30 platform 200.

Docket No. AUS9-2000-0220-US1

Partitioned hardware **230** includes a plurality of processors **232-238**, a plurality of system memory units **240-246**, a plurality of input/output (I/O) adapters **248-262**, and a storage unit **270**. Each of the processors

5 **242-248**, memory units **240-246**, and I/O adapters **248-262** may be assigned to one of multiple partitions within logically partitioned platform **200**, each of which corresponds to one of operating systems **202-208**.

Hypervisor **210**, implemented as firmware, creates and

10 enforces the partitioning of logically partitioned platform **200**. Firmware is software stored in a memory chip that holds its content without electrical power, such as, for example, read-only memory (ROM), programmable ROM (PROM), erasable programmable ROM

15 (EPROM), electrically erasable programmable ROM (EEPROM), and non-volatile random access memory (non-volatile RAM).

Virtual address translation hardware **280** provides a mechanism for translating one of OSs' **202-208** virtual memory address pages for a resource to a physical

20 hardware address corresponding to that resource. Virtual memory is a method of simulating more memory than actually exists, allowing platform **200** to run larger software programs or more programs concurrently. Virtual memory breaks up the software program into small

25 segments, called pages, and brings as many pages into memory **240-246** that fit into a reserved area for that software program. When additional pages are required, the virtual memory makes room for them by swapping pages currently in memory but no longer needed to disk storage

30 **270** or some other input/output device through one of I/O

Docket No. AUS9-2000-0220-US1

adapters **248-260**, thus freeing up memory for the additional pages. The virtual address translation hardware **280** keeps track of pages that have been modified, so that the modified pages can be retrieved
5 when needed again.

In the prior art, the platform's virtual address translation hardware **280** was under the direct control of each OS **202-208** within the platform **200**. However, such direct control allowed for the possibility that one of
10 OSs **202-208** might read or corrupt data in a resource allocated to a different one of OSs **202-208** than the requesting one of OSs **202-208**. Thus, the present invention removes the virtual address translation hardware **280** from the direct control of the OSs **202-208**.
15 Instead, hypervisor **210** controls platform's **200** virtual address translation hardware **280**.

When one of OSs **202-208** needs access to a specific resource, the particular one of OSs **202-208** must establish a virtual to real mapping for the input/output
20 resource. However, instead of directly modifying the virtual address translation page frame table, the OS **202-208** makes a request to hypervisor **210** to access the resource. Hypervisor **210** checks that the real resource, to which one of OSs **202-208** has requested access, has
25 been allocated to the requesting one of OSs **202-208** before hypervisor **210** completes the mapping. If the real resource has been allocated to the one of OSs **202-208** requesting the mapping, hypervisor **210** then performs the final stage of the mapping.

30 The single hardware pointer to the virtual address

Docket No. AUS9-2000-0220-US1

translation page frame table in each of processors
232-238 is made unmodifiable by any of OSs **202-208**. Each
processor **232-238** contains a register that points to its
page table, such that the processor's virtual memory
5 translation hardware is aware of where to find the
virtual to physical address mappings. This is standard
art for memory address translation hardware. In the
depicted embodiment, each OS image **202-208** has its own
page table such that each OS image **202-208** can have its
10 own virtual address 1, 1000, 1000000, etc. Furthermore,
the page table pointer register of each of processors
232-238 assigned to a particular one of OS images 202-208
contains the same value (i.e. the starting physical
address of the page table in platform memory). Each of
15 OSs **202-208** contains an instruction that allows the OS to
make request of the hypervisor **210** to perform the virtual
address resource mapping.

Those of ordinary skill in the art will appreciate
that the hardware and software depicted in **Figure 2** may
20 vary. For example, more or fewer processors and/or more
or fewer operating system images may be used than those
depicted in **Figure 2**. The depicted example is not meant
to imply architectural limitations with respect to the
present invention.

25 With reference now to **Figure 3**, a block diagram
illustrating an exemplary system for mediating virtual
address translation is depicted in accordance with the
present invention. Mediated address translation system
300 includes multiple operating systems (OSs) **302-308**, a
30 hypervisor **310**, an allocation table **380**, a page frame

Furthermore, there is no requirement that a page
30 frames **361-371** in physical resources **360** be allocated

Docket No. AUS9-2000-0220-US1

such that each of the page frames **361-371** allocated to a particular partition be grouped consecutively. Thus, the page frames **361-371** may be assigned in a disjointed fashion such that, for example, page frame **361** is

5 assigned to OS **302**, page frame **363** is assigned to OS **304**, page frame **371** is assigned to OS **308**, page frame **362** is assigned to OS **302**, page frame **369** is assigned to OS **306**, and so on. Therefore, the present invention allows for a very fine grain allocation of resources to the various

10 logical partitions. Furthermore, the allocations of physical resources are easily modified by changes to hypervisor address checking tables, such as, for example, the allocation table **380**, without the need for hardware reconfiguration as some prior art methods required.

15 Each of OSs **302-308**, when access to a physical resource **350** is desired, sends a request to hypervisor **310** to map a virtual address into a physical address. Hypervisor **310**, after receiving the request, consults an allocation table to ensure that the physical resource

20 requested by the OS **302-308** has been allocated to that particular requesting OS **302-308**. If hypervisor **310** determines that the requested resource has not been allocated to the requesting OS **302-308**, then hypervisor **310** returns an error message to the requesting OS **302-308**

25 indicating that the request has been denied. Thus, hypervisor **310** prevents an OS **302-308** in one partition from fetching or corrupting data in a resource allocated to an OS **302-308** in a different partition, thereby maintaining the integrity of the logically partitioned

30 data processing system.

Docket No. AUS9-2000-0220-US1

If the requested resource has been allocated to the requesting OS **302-308**, then hypervisor **310** modifies page frame table **320**, if necessary, such that the OS's **302-308** virtual address is mapped to the corresponding physical address of the requested resource. The various OSs **302-308** are prevented from modifying page frame table **320**, thus further ensuring that the logical partitions within the data processing system are maintained.

Other prior art methods incorporated the virtual address translation hardware into their firmware. However, these other prior methods incorporated the entire virtual translation function into their firmware rather than merely allowing the firmware to modify the page frame table after consulting an allocation table. Thus, these other prior art methods forced significantly greater changes in a standard UNIX operating system and they introduced other mechanisms for partitioning the I/O that greatly complicated their design. Thus, the present invention is much simpler to incorporate into existing OSs and simpler to implement in firmware.

With reference now to **Figure 4**, a flowchart illustrating an exemplary process in a firmware component, such as, for example, hypervisor **210** in **Figure 2**, for mediating address translation between various operating systems in a logically partitioned data processing system and physical resources is depicted in accordance with the present invention. To begin, the firmware component receives a request from one of multiple operating systems within a logically partitioned platform to access (read or write to) a physical resource, such as, for example, an I/O device connected

Docket No. AUS9-2000-0220-US1

to I/O adapter **248** in **Figure 2** (step **402**). The firmware component consults an allocation table (step **404**) to determine whether the resource to which the OS requested access has been allocated to the requesting OS (step
5 **406**).

After consulting the allocation table, if the firmware component makes the determination that the requested resource has not been allocated to the requesting OS, then an error message is returned to the
10 requesting OS indicating that access to such resource is denied (step **408**). In such case, the firmware component does not make any modifications to the page frame table, which maps the OSs virtual addresses to the physical addresses of the corresponding resources. If the
15 firmware component makes the determination that the requested resource has been allocated to the requesting OS, then the firmware component modifies the page frame table such that the OSs virtual address is mapped to the resources physical address (step **410**).

20 It is important to note that while the present invention has been described in the context of a fully functioning data processing system, those of ordinary skill in the art will appreciate that the processes of the present invention are capable of being distributed in
25 the form of a computer readable medium of instructions and a variety of forms and that the present invention applies equally regardless of the particular type of signal bearing media actually used to carry out the distribution. Examples of computer readable media
30 include recordable-type media such as a floppy disc, a hard disk drive, a RAM, and CD-ROMs and transmission-type

media such as digital and analog communications links.

5
10

Docket No. AUS9-2000-0220-US1

CLAIMS:

What is claimed is:

1. A method for mediating address translation in a
5 logically partitioned data processing system having a set of logical partitions with an operating system assigned to each logical partition within the set of logical partitions, the method comprising:
receiving from an operating system within a logical
10 partition from the set of logical partitions a request to access a physical resource;
responsive to a determination that the physical resource has been allocated to the logical partition, selectively modifying an address translation table to
15 allow access to the physical resource by the operating system.
2. The method as recited in claim 1, further comprising:
20 responsive to a determination that the physical resource is allocated to a different logical partition in the set of logical partitions, refraining from modifying the address translation table.
- 25 3. The method as recited in claim 2, further comprising:
sending a message to the operating system indicating that the request is denied.
- 30 4. The method as recited in claim 1, wherein the address translation table comprises a table of virtual

Docket No. AUS9-2000-0220-US1

addresses with corresponding physical addresses, wherein the virtual addresses are addresses utilized by the operating system and the physical addresses are addresses corresponding to the physical location of resources
5 within the logically partitioned data processing system.

5. The method as recited in claim 4, wherein the physical addresses are allocated to various ones of multiple logical partitions in a disjoint fashion.
10

6. The method as recited in claim 4, wherein consecutive virtual addresses need not correspond to consecutive physical addresses.

7. A computer program product in a computer readable media for use in a logically partitioned data processing system for mediating address translation in a logically partitioned data processing system having a set of logical partitions with an operating system assigned to each logical partition in the set of logical partitions,
15 20 the computer program product comprising:

first instructions for receiving from an operating system within a logical partition from the set of logical partitions a request to access a physical resource;

25 second instructions, responsive to a determination that the physical resource has been allocated to the logical partition, for selectively modifying an address translation table to allow access to the physical resource by the operating system.

30

8. The computer program product as recited in claim 7,

further comprising:

5

further comprising:

10

15

20

25

30

Docket No. AUS9-2000-0220-US1

logically partitioned data processing system having a set of logical partitions with an operating system assigned to each logical partition in the set of logical partitions, the system comprising:

5 first means for receiving from an operating system within a logical partition from the set of logical partitions a request to access a physical resource;

 second means, responsive to a determination that the physical resource has been allocated to the logical
10 partition, for selectively modifying an address translation table to allow access to the physical resource by the operating system.

14. The system as recited in claim 13, further
15 comprising:

 third means, responsive to a determination that the physical resource is allocated to a different logical partition in the set of logical partitions, for
20 refraining from modifying the address translation table.

15. The system as recited in claim 14, further
 comprising:

 fourth means for sending a message to the operating system indicating that the request is denied.

25

16. The system as recited in claim 13, wherein the address translation table comprises a table of virtual addresses with corresponding physical addresses, wherein the virtual addresses are addresses utilized by the
30 operating system and the physical addresses are addresses corresponding to the physical location of resources

Docket No. AUS9-2000-0220-US1

within the logically partitioned data processing system.

17. The system as recited in claim 16, wherein the
physical addresses are allocated to various ones of
5 multiple logical partitions in a disjoint fashion.

18. The system as recited in claim 16, wherein
consecutive virtual addresses need not correspond to
consecutive physical addresses.

10

19. A logically partitioned data processing system,
comprising:

a plurality of operating systems executing within
the logically partitioned data processing system, each of
15 the plurality of operating systems assigned to one of a
plurality of logical partitions;

a plurality of physical resources, each assigned to
one of the plurality of logical partitions; and

a mediating component for providing address
20 translation between each of a plurality of virtual
addresses belonging to various ones of the plurality of
operating systems and a corresponding one of a plurality
of physical addresses belonging to various ones of the
plurality of physical resources; wherein

25 the mediating component determines whether a
requested resource has been allocated to a requesting one
of the plurality of operating systems before mapping a
one of the plurality of virtual addresses to a one of the
plurality of physical addresses belonging to the
30 requested resource; and

if it is determined that the requested resource is

not allocated to the logical partition to which the requesting one of the plurality of operating systems is allocated, the mediating component refrains from mapping the one of the plurality of virtual addresses to the one of the plurality of physical addresses belonging to the requested resource.

5
10
15
20

LOGICAL PARTITIONING VIA HYPERVISOR MEDIATED ADDRESS TRANSLATION

A method, system, and apparatus for mediating address translation in a logically partitioned data processing system is provided. In one embodiment, a firmware component receives from an operating system within a logical partition a request to access a physical resource. The firmware component, responsive to a determination that the physical resource has been allocated to the logical partition, modifies an address translation table, if necessary, to allow access to the physical resource by the operating system. The operating system is prevented from directly modifying the address translation table, thus preventing potential interference between operating systems within the logically partitioned data processing system.

20

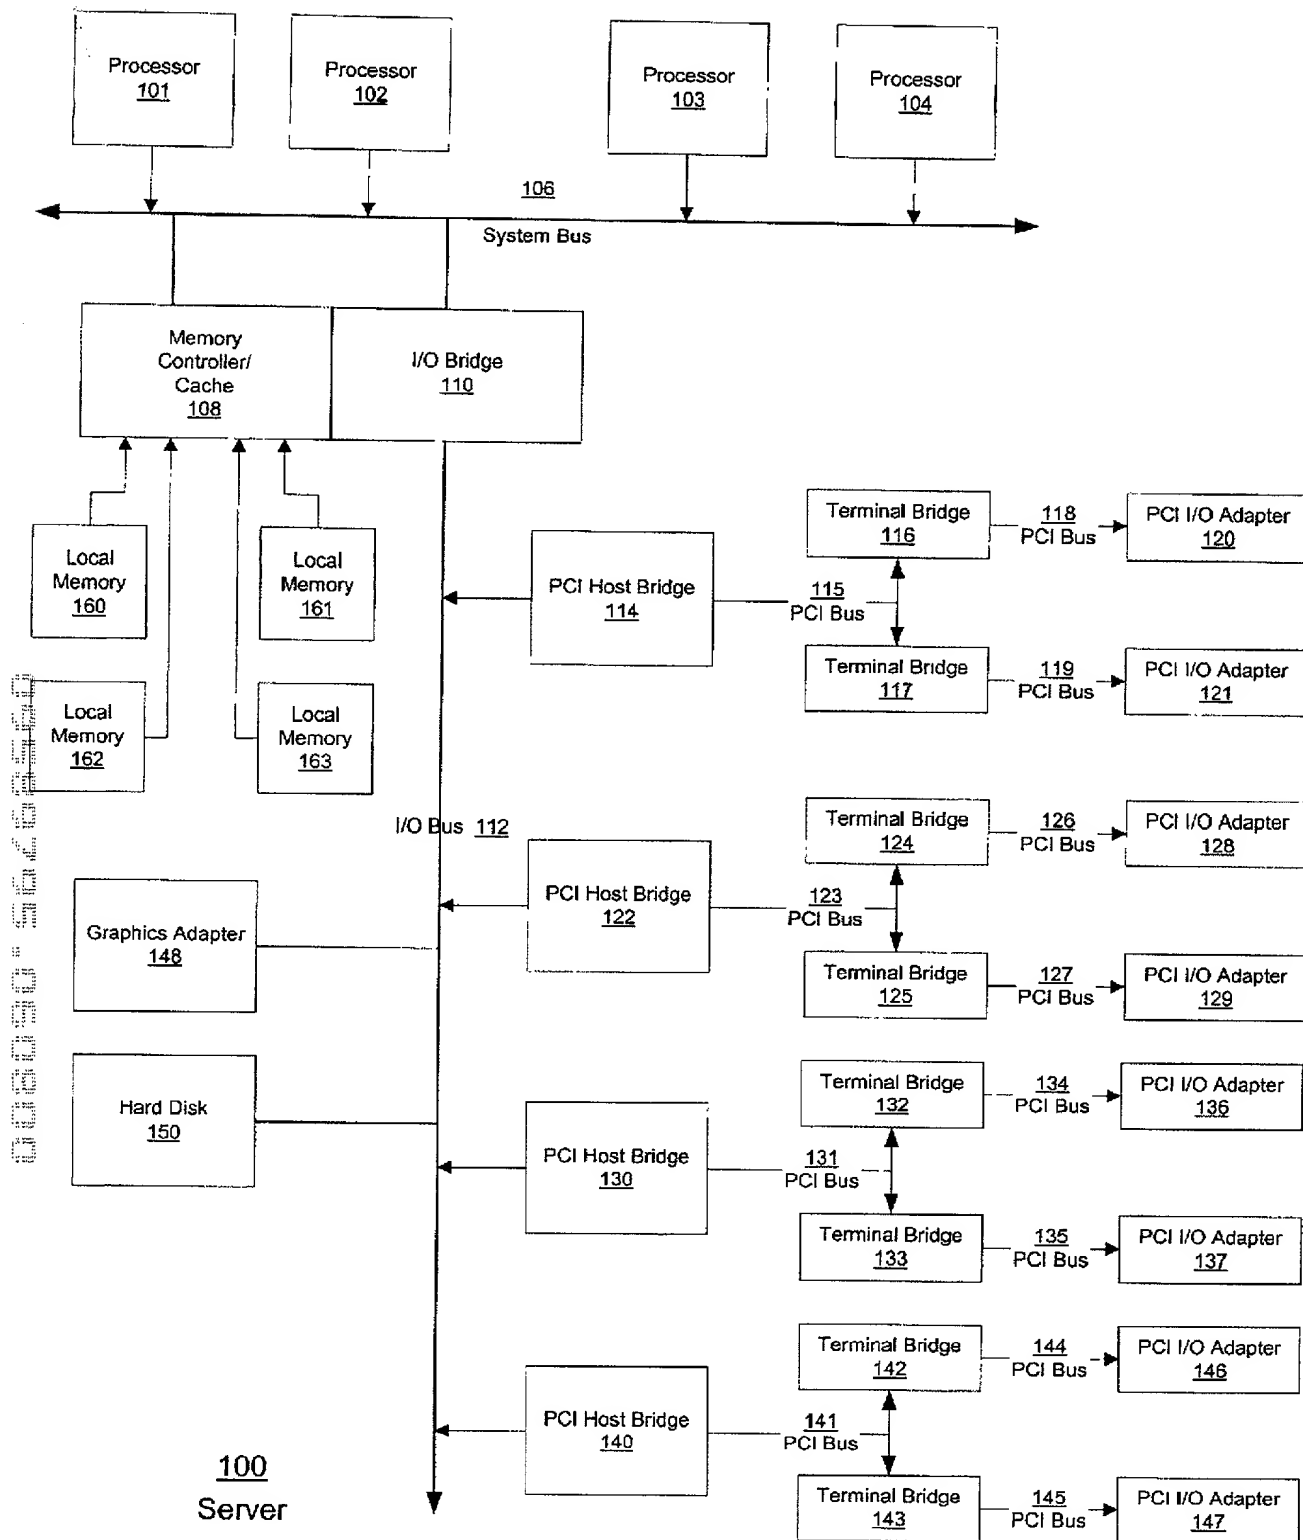


Figure 1

AUS9-2000-0220-US1

OS
202

OS
204

OS
206

OS
208

Hypervisor
210

Processor
232

Processor
234

Processor
236

Processor
238

Storage
270

Memory
240

Memory
242

Memory
244

Memory
246

I/O Adapter
248

I/O Adapter
250

I/O Adapter
252

I/O Adapter
254

I/O Adapter
256

I/O Adapter
258

I/O Adapter
260

I/O Adapter
262

Partitioned Hardware
230

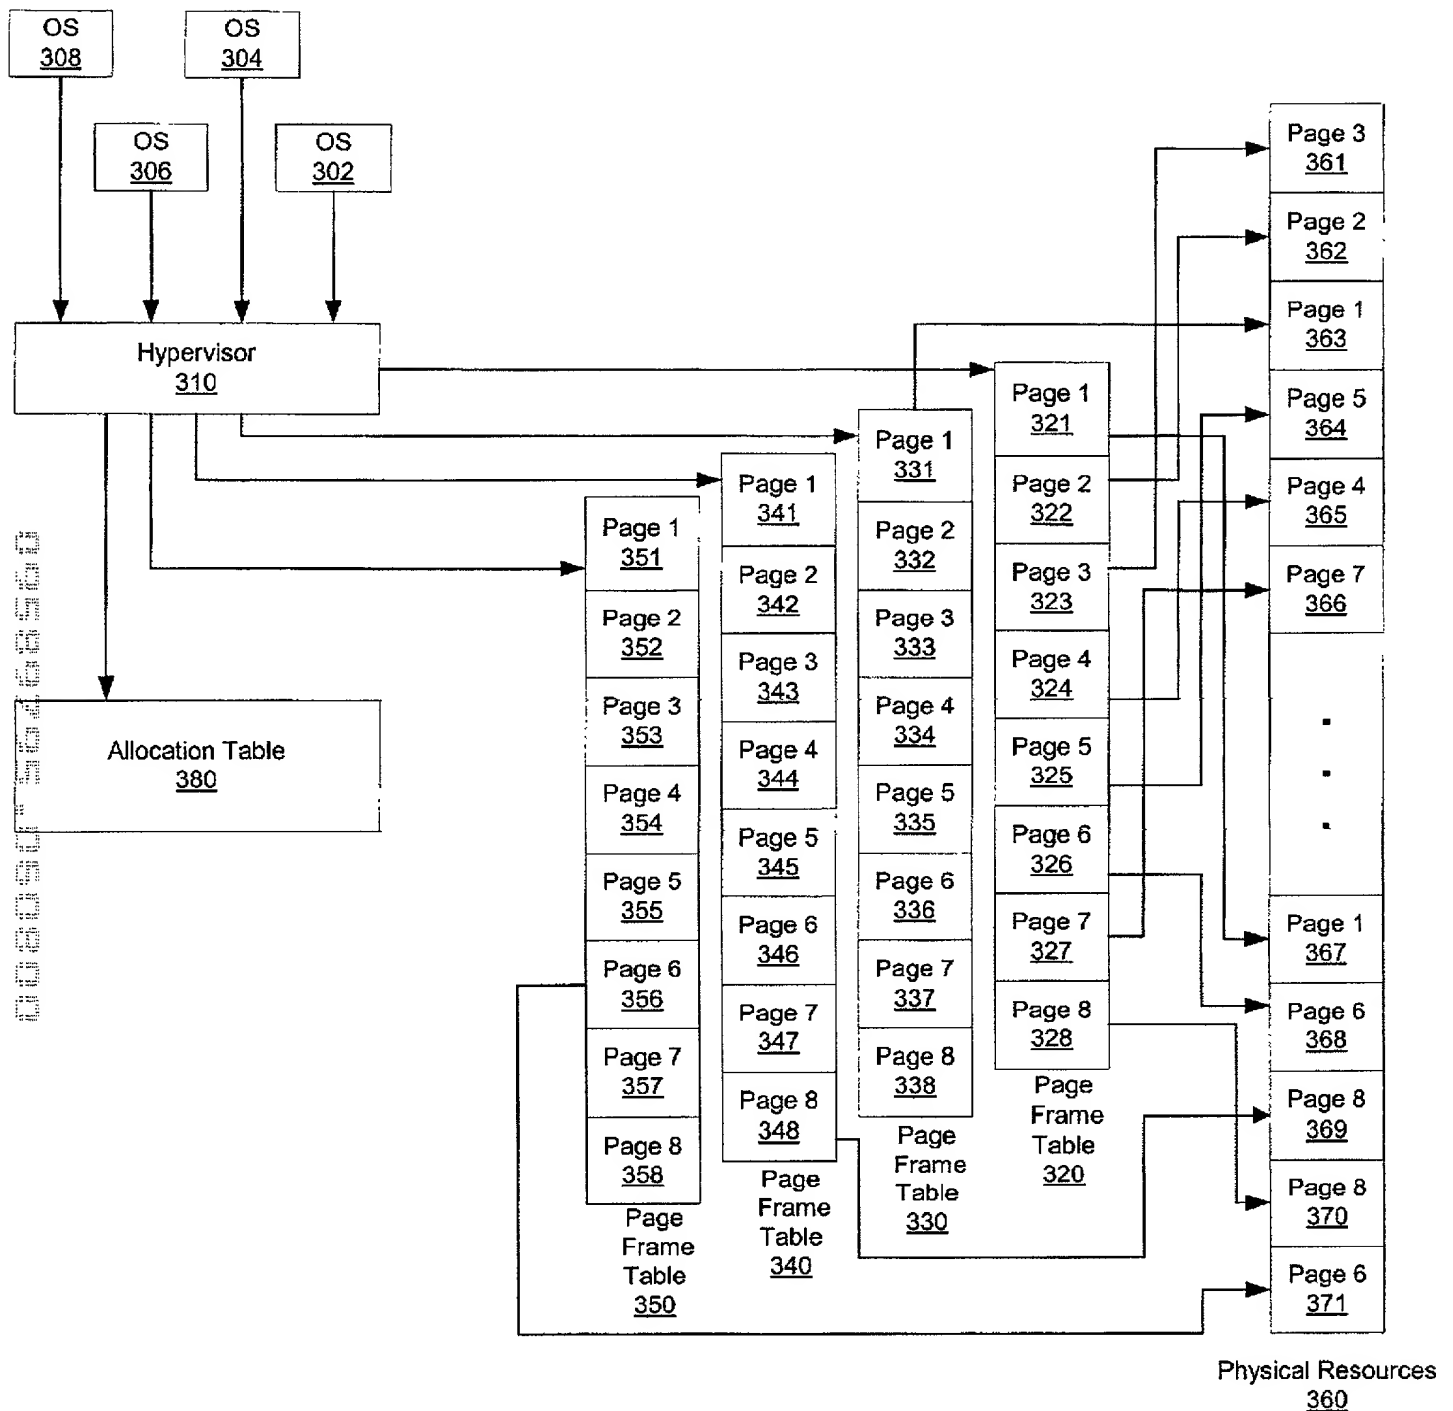
Virtual Address Translation Hardware
280

200
Logically Partitioned Platform

AUS9-2000-0220-US1

Figure 3

AUS9-2000-0220-US1



```

graph TD
    Start([Start]) --> 402[Receive request from operating system to access resource.  
402]
    402 --> 404[Consult allocation table to determine whether resource is  
allocated to requesting OS.  
404]
    404 --> 406{Is  
resource allocated to  
requesting OS?  
406}
    406 -- NO --> 408[Return error message to  
requesting OS.  
408]
    406 -- YES --> 410[Map OSs virtual address to resources physical address in page  
frame table.  
410]
    408 --> Stop([Stop])
    410 --> Stop
  
```

AUS9-2000-0220-US1

**DECLARATION AND POWER OF ATTORNEY FOR
PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

LOGICAL PARTITIONING VIA HYPERVISOR MEDIATED ADDRESS TRANSLATION

the specification of which (check one)

☒ is attached hereto.

— was filed on _____
as Application Serial No. _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s): Priority Claimed

(Number) (Country) (Day/Month/Year) ___ Yes___ No

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #) (Filing Date) (Status)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; Thomas E. Tyson, Reg. No. 28,543; James H. Barksdale, Jr., Reg. No. 24,091; Casimer K. Salys, Reg. No. 28,900; Robert M. Carwell, Reg. No. 28,499; Douglas H. Lefevre, Reg. No. 26,193; Jeffrey S. LaBaw, Reg. No. 31,633; David A. Mims, Jr., Reg. 32,708; Volel Emile, Reg. No. 39,969; Anthony V. England, Reg. No. 35,129; Leslie A. Van Leeuwen, Reg. No. 42,196; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Marilyn S. Dawkins, Reg. No. 31,140; Mark E. McBurney, Reg. No. 33,114; Duke W. Yee, Reg. No. 34,285; Colin P. Cahoon, Reg. No. 38,836; Joseph R. Burwell, Reg. No. 44,468; Rudolph J. Buchel, Reg. No. 43,448; and Stephen R. Loe, Reg. No. 43,757, Stephen J. Walder, Reg. No. 41,534.

Send correspondence to: Duke W. Yee, Carstens, Yee & Cahoon, LLP, P.O. Box 802334, Dallas, Texas 75380 and direct all telephone calls to Duke W. Yee, (972) 367-2001

FULL NAME OF SOLE OR FIRST INVENTOR: RICHARD LOUIS ARNDT

INVENTORS SIGNATURE: Richard Louis Arndt DATE: 7 June 2000

RESIDENCE: 1607 BARN SWALLOW DRIVE
AUSTIN, TEXAS 78746

CITIZENSHIP: UNITED STATES

POST OFFICE ADDRESS: SAME AS ABOVE